arXiv:1703.06277v1 [stat.ME] 18 Mar 2017

Unsupervised Learning of Mixture Regression Models for Longitudinal Data

Peirong Xu[1], Heng Peng[2], Tao Huang[3*]

[1] College of Mathematics and Sciences, Shanghai Normal University, Shanghai, China

[2] Department of Mathematics, Hong Kong Baptist University, Hong Kong, China

[3] School of Statistics and Management, Shanghai University of Finance and Economics, Shanghai, China

Abstract: This paper is concerned with learning of mixture regression models for individuals that are measured repeatedly. The adjective "unsupervised" implies that the number of mixing components is unknown and has to be determined, ideally by data driven tools. For this purpose, a novel penalized method is proposed to simultaneously select the number of mixing components and to estimate the mixing proportions and unknown parameters in the models. The proposed method is capable of handling both continuous and discrete responses by only requiring the first two moment conditions of the model distribution. It is shown to be consistent in both selecting the number of components and estimating the mixing proportions and unknown regression parameters. Further, a modified EM algorithm is developed to seamlessly integrate model selection and estimation. Simulation studies are conducted to evaluate the finite sample performance of the proposed procedure. And it is further illustrated via an analysis of a primary biliary cirrhosis data set.

Key words: Unsupervised learning, Model selection, Longitudinal data analysis, Quasi-likelihood, EM algorithm.

*__Corresponding Author:__ Tao Huang, Associate Professor, School of Statistics and Management, Shanghai University of Finance and Economics, Shanghai, China, 200433. (*E-mail:huang.tao@mail.shufe.edu.cn*).

1 Introduction

In many medical studies, the marker of disease progression and a variety of characteristics are routinely measured during the patients' follow-up visit to decide on future treatment actions. Consider a motivating Mayo Clinic trial with primary biliary cirrhosis (PBC), wherein a number of serological, clinical and histological parameters were recorded for each of 312 patients from 1974 to 1984. This longitudinal study had a median follow-up time of 6.3 years as some patients missed their appointments due to worsening medical condition of some labs. It is known that PBC is a fatal chronic cholesteric liver disease, which is characterized histopathologically by portal inflammation and immune-mediated destruction of the intrahepatic bile ducts. It can be divided into four histologic stages, but with nonuniformly affected liver. The diagnosis of PBC is important for the medical treatment with Ursodiol has been shown to halt disease progression and extend survival free of liver transplantation. Therefore, one goal of the study was the investigation of the serum bilirunbin level, an important marker of PBC progression, in relation to the time and to potential clinical and histological covariates. Another issue that should be accounted for is the unobservable heterogeneity between subjects that may not be explained by the covariates. The changes in inflammation and bile ducts occur at different rates and with varying degrees of severity in different patients, so the heterogeneous patients could potentially belong to different latent groups. To address these problems, there is a demand for mixture regression modeling for subjects on the basis of longitudinal measurements.

There are various research works on mixture regression models for longitudinal outcome data, particularly in the context of model-based probabilistic clustering (Fraely and Raftery, 2002). For example, De la Cruz-Mesía et. al. (2008) proposed a mixture of non-linear hierarchical models with Gaussian subdistributions; McNicholas and Murphy (2010) extended the Gaussian mixture models with Cholesky-decomposed group covariance structure; Komárek and Komárková (2013) introduced a generalized linear mixed model for components' densities under the Gaussian mixture framework; Heinzl

and Tutz (2013) considered linear mixed models with approximate Dirichlet process mixtures. Other relevant work includes Celeux et. al. (2005), Booth et. al. (2008), Pickles and Croudace (2010), Maroutti (2011), Erosheva et. al. (2014) and some of the references therein. Compared with heuristic methods such as the k-means method (Genolini and Falissard, 2010), issues like the selection of the number of clusters (or components) can be addressed in a principled way. However, most of them assume a parametric mixture distribution, which may be too restrictive and invalid in practice when the true data-generating mechanism indicates otherwise.

A key concern for the performance of mixture modeling is the selection of the number of components. A mixture with too many components may overfit the data and result in poor interpretations. Many statistical methods have been proposed in the past few decades by using the information criteria. For example, see Leroux (1992), Roeder and Wasserman (1997), Hennig(2004), De la Cruz-Mesía et al. (2008) and many others. However, these methods are all based on the complete model search algorithm, which result in heavy computation burden. To improve the computational efficiency, data-driven procedures are much more preferred. Recently, Chen and Khalili (2008) used the SCAD penalty (Fan and Li, 2001) to penalize the difference of location parameters for mixtures of univariate location distributions; Komárek and Lesaffre (2008) suggested to penalize the reparameterized mixture weights in the generalized mixed model with Gaussian mixtures; Heinzl and Tutz (2014) constructed a group fused lasso penalty in linear-mixed models; Huang et. al. (2016) proposed a penalized likelihood method in finite Gaussian mixture models. Most of them are developed for independent data or based on the full likelihood. However, the full likelihood is often difficult to specify in formulating a mixture model for longitudinal data, particularly for correlated discrete data.

In this paper, we propose a new penalized method for mixture regression models to deal with the above mentioned problems simultaneously. This would be the first attempt to handle both balanced and unbalanced longitudinal data that only requires the first two moment conditions of the model distribution. By penalizing the logarithm of mixing

proportions, our approach can simultaneously select the number of mixing components and estimate the mixing proportions and unknown parameters in the semiparametric mixture regression model. The number of components can be consistently selected. And given the number of components, the estimators of mixing proportions and regression parameters can be root-n consistent and asymptotically normal. By taking account of the within-component dispersion, we further develop a modified EM algorithm to improve the classification accuracy. Simulation results and the application to the motivating PBC data demonstrate the feasibility and effectiveness of the proposed method.

The rest of the paper is organized as follows. In Section 2, we introduce a new penalized method for learning semiparametric mixture regression models with longitudinal data. Section 3 presents the corresponding theoretical properties and Section 4 provides a modified EM algorithm for implementation. In Section 5, we assess the finite sample performance of the proposed method via simulation studies. We apply the proposed method to the PBC data in Section 6, and conclude the paper with Section 7. All technical proofs are provided in Appendix.

2 Learning semiparametric mixture of regressions

2.1 Model specification

In a longitudinal study, suppose Y_{ij} is the response variable measured at the jth time point for the ith subject, and X_{ij} is the corresponding $p \times 1$ vector of covariates, $i = 1, \ldots, n$, $j = 1, \ldots, m_i$. Let $Y_i = (Y_{i1}, \ldots, Y_{im_i})^T$ and $X_i = (X_{i1}, \ldots, X_{im_i})^T$. In general, the observations for different subjects are independent, but they may be correlated within the same subject. We assume that the observations of each subject belong to one of K classes (components) and $u_i \in \{1, \ldots, K\}$ is the corresponding latent class variable. Conditioning on X_i, assume that u_i has a discrete distribution $\mathrm{P}(u_i = k \mid X_i) = \pi_k$, where $\pi_k, k = 1, \ldots, K$, are the positive mixing proportions satisfying $\sum_{k=1}^{K} \pi_k = 1$.

Given $u_i = k$ and X_{ij}, suppose the conditional mean of Y_{ij} is

$$\mu_{ijk} \equiv \mathrm{E}(Y_{ij} \mid X_{ij}, u_i = k) = g(X_{ij}^T \beta_k), \tag{2.1}$$

where g is a known link function, and β_k is a $p-$dimensional unknown parameter vector. The corresponding conditional variance of Y_{ij} is given by

$$\sigma_{ijk}^2 \equiv \mathrm{var}(Y_{ij} \mid X_{ij}, u_i = k) = \phi_k V(\mu_{ijk}), \tag{2.2}$$

where V is a known positive function and ϕ_k is a unknown dispersion parameter. In other words, conditioning on X_{ij}, the response variable Y_{ij} follows a mixture distribution

$$Y_{ij} \mid X_{ij} \sim \sum_{k=1}^{K} \pi_k f_k(Y_{ij} \mid X_{ij}^T \beta_k, \phi_k),$$

where $f_k(Y_{ij} \mid X_{ij}^T \beta_k, \phi_k)$'s are the components' distributions. To avoid identifiability issues, we assume that K is the smallest integer such that $\pi_k > 0$ for $k = 1, \cdots, K$, and $(\beta_a, \phi_a) \neq (\beta_b, \phi_b)$ for $1 \leq a < b \leq K$. Denote $\theta = (\beta_1^T, \ldots, \beta_K^T, \phi^T, \pi^T)^T$ with $\beta_k = (\beta_{k1}, \ldots, \beta_{kp})^T$ and $\pi = (\pi_1, \ldots, \pi_{K-1})^T$, and $\phi = (\phi_1, \ldots, \phi_K)^T$.

Under the working independence correlation, the (log) quasi-likelihood of the K-component marginal mixture regression model is

$$Q(\theta) = \sum_{i=1}^{n} \log \left[\sum_{k=1}^{K} \pi_k \exp \left\{ \sum_{j=1}^{m_i} q(g(X_{ij}^T \beta_k); Y_{ij}) \right\} \right], \tag{2.3}$$

where function $q(\mu; y)$ (McCullagh and Nelder, 1989) satisfies $\frac{\partial q(\mu;y)}{\partial \mu} = \frac{y-\mu}{V(\mu)}$. It is known that, for a generalized linear model with independent data, the quasi-likelihood estimator of the regression coefficient has the same asymptotic properties as the maximum likelihood estimator. While for longitudinal data, it is equivalent to the GEE estimator (Liang and Zeger, 1986), which is consistent even when the working correlation structure is misspecified. Therefore, estimation consistency is expected to hold for the K-component marginal mixture regression model (2.1)-(2.2), and this will be validated in Section 3.

2.2 Penalized quasi-likelihood method

For a fixed number of K components, we can maximize the quasi-likelihood function (2.3) by an expectation-maximization (EM) algorithm, which in the E-step computes the

posterior probability of the class memberships and in the M-step estimates the mixing proportions and unknown parameters. However, in practice, the number of components is usually unknown and needs to be inferred from the data itself.

For the proposed marginal mixture regression model, the selection of the number of mixing components can be viewed as a model selection problem. Various conventional methods have been proposed based on the likelihood function and some information theoretic criteria. In particular, the Bayesian information criterion (BIC; Schwarz, 1978) is recommended as a useful tool for selecting the number of components (Dasgupta and Raftery, 1998; Fraley and Rafetery, 2002). Therefore, a natural idea is to propose a BIC-type criterion for selecting the number of mixing components, where the likelihood function is replaced by the quasi-likelihood function (2.3). But our simulation experience shows that it couldn't perform as well as the traditional BIC, since (2.3) is no longer a joint density with integral equals to one.

To avoid calculating the normalizing constant, the penalization technique is preferred. By (2.3), intuitively, the kth component would be eliminated if $\pi_k = 0$. But in implementation of (2.3), the quasi-likelihood function for the complete data (u_{ik}, Y_i, X_i) involves $\log \pi_k$ rather than π_k, where u_{ik} denotes the indicator of whether ith subject belongs to the kth component (see (4.1) defined in Section 4 for details). Therefore, it is natural to penalize the logarithm of mixing proportions $\log \pi_k$, $k = 1, \ldots, K$. Moreover, note that the gradient of $\log \pi_k$ increases very fast when π_k is close to zero, and it would dominate the gradient of nonzero $\pi_l > 0$. Consequently, the popular L_q types of penalties may not able to set insignificant π_k to zero. In the spirit of penalization in Huang et al. (2016), we propose the following penalized quasi-likelihood function

$$Q_{\mathrm{P}}(\theta) = Q(\theta) - n\lambda \sum_{k=1}^{K} \{\log(\epsilon + \pi_k) - \log(\epsilon)\}, \tag{2.4}$$

where λ is a tuning parameter and ϵ is a very small positive constant. Note that $\log(\epsilon + \pi_k) - \log(\epsilon)$ is an increasing function of π_k and is shrunk to zero as the mixing proportion π_k goes to zero. Therefore, the proposed method (2.4) can simultaneously

determine the number of mixture components and estimate mixing proportions and unknown parameters.

REMARK 1. *The small constant ϵ is introduced to ensure the continuity of the objective function when some of mixing proportions are shrunk continuously to zero.*

REMARK 2. *The penalty $n\lambda\sum_{k=1}^{K}\{\log(\epsilon+\pi_k)-\log(\epsilon)\}$ in (2.4) would over penalize large π_k and result in a biased estimator. A more general but slightly more complicated approach is to use $n\sum_{k=1}^{K}\{\log(\epsilon+p_\lambda(\pi_k))-\log(\epsilon)\}$, where $p_\lambda(\cdot)$ is a penalty function that gives estimators with sparsity, unbiasedness and continuity as discussed in Fan and Li (2001).*

3 Asymptotic properties

In this section, we first study the asymptotic property of the maximum quasi-likelihood estimator $\widehat{\theta}$ of (2.3) given the number of mixing components. And then, we establish the model selection consistency of the proposed method (2.4) for the general semiparametric marginal mixture regression model (2.1)-(2.2).

For a fixed number of K components, denote the true value of parameter vector by θ_0. The components of θ_0 are denoted with a subscript, such as π_{0k}. We assume that the number of subjects n increases to infinity, while the number of observations $\{m_i\}$ is a bounded sequence of positive integers. Let

$$\Psi(\theta;Y_i|X_i)=\sum_{k=1}^{K}\pi_k\exp\left\{\sum_{j=1}^{m_i}q(g(X_{ij}^T\beta_k);Y_{ij})\right\} \tag{3.1}$$

and $\psi(\theta;Y_i\mid X_i)=\log(\Psi(\theta;Y_i\mid X_i))$.

We assume the following regularity conditions to derive the asymptotic properties.

C1 The function $g(\cdot)$ has two bounded and continuous derivatives.

C2 The random variables X_{ij}'s are bounded on the compact support $\mathcal{A}$ uniformly. For $\theta\in\Omega$, the density function of $X_{ij}^T\beta_k$ is positive and satisfies Lipschitz condition of order 1 on $\mathcal{U}_k=\{u=X_{ij}^T\beta_k : X_{ij}\in\mathcal{A}, i=1,\ldots,n, j=1,\ldots,m_i\}$, $k=1,\ldots,K$.

C3 For each $\theta \in \Omega$, $\psi(\theta; Y_i \mid X_i)$ admits third order partial derivatives with respect to θ. And there exist functions $M_l(X_i, Y_i)$, $l = 1, 2, 3$ such that for θ in a neighborhood of θ_0, $|\partial\psi(\theta; Y_i \mid X_i)/\partial\theta_j| \leq M_1(X_i, Y_i)$, $|\partial^2\psi(\theta; Y_i \mid X_i)/\partial\theta_j\partial\theta_k| \leq M_2(X_i, Y_i)$, and $|\partial^3\psi(\theta_0; Y_i \mid X_i)/\partial\theta_j\partial\theta_k\partial\theta_l| \leq M_3(X_i, Y_i)$ with $E\{M_l(X_i, Y_i)\} < \infty$, for all $i = 1, \ldots, n$.

C4 Let $A = \text{var}\{\partial\psi(\theta_0; Y_i \mid X_i)/\partial\theta\}$. The second derivative matrix $B = \text{E}\{-\partial^2\psi(\theta_0; Y_i \mid X_i)/\partial\theta\partial\theta^T\}$ is positive definite.

Conditions C1-C2 are typical assumptions in the estimation literature, which are also found in Xu and Zhu (2012) and Xu et. al. (2016). Conditions C3 and C4 are mild conditions in the literature of mixture models, which are used for the proof of asymptotic normality.

THEOREM 1. *Under conditions C1-C4, the maximum quasi-likelihood estimator $\widehat{\theta}$ of (2.3) given the number of components has the asymptotic normality*

$$\sqrt{n}(\widehat{\theta} - \theta_0) \xrightarrow{L} N(0, B^{-1}AB^{-1}).$$

Next, we study the model selection consistency of the proposed method (2.4) for the marginal mixture regression model (2.1)-(2.2). We assume that there are K_0 mixture components, $K_0 \leq K$ with $\pi_l = 0$, for $l = 1, \ldots, K - K_0$ and $\pi_l = \pi_{0k}$ for $l = K - K_0 + 1, \ldots, K$, $k = 1, \ldots, K_0$. In the spirit of locally conic parametrization (Dacunha-Castelle and Gassiat, 1997), define $\pi_l = \lambda_l\eta$, $l = 1, \ldots, K - K_0$ and $\pi_l = \pi_{0k} + \rho_k\eta$, $l = K - K_0 + 1, \ldots, K$, $k = 1, \ldots, K_0$. Then, the function (3.1) can be rewritten as

$$\Psi(\eta, \gamma; Y_i \mid X_i) = \sum_{l=1}^{K-K_0} \lambda_l\eta f(\beta_l; Y_i \mid X_i) + \sum_{k=1}^{K_0} (\pi_{0k} + \rho_k\eta) f(\beta_{0k} + \eta\delta_k; Y_i \mid X_i),$$

where $f(\beta_l; Y_i \mid X_i) = \exp\left\{\sum_{j=1}^{m_i} q(g(X_{ij}^T\beta_l); Y_{ij})\right\}$ and $\gamma = (\lambda_1, \ldots, \lambda_{K-K_0}, \rho_1, \ldots, \rho_{K_0}, \beta_1^T, \ldots, \beta_{K-K_0}^T, \delta_1^T, \ldots, \delta_{K_0}^T)^T$ with restrictions $\lambda_l \geq 0, \beta_l \in R^p, l = 1, \ldots, K - K_0$, $\delta_k \in R^p, \rho_k \in R, k = 1, \ldots, K_0$, $\sum_{l=1}^{K-K_0} \lambda_l + \sum_{k=1}^{K_0} \rho_k = 0$ and $\sum_{l=1}^{K-K_0} \lambda_l^2 + \sum_{k=1}^{K_0} \rho_k^2 + \sum_{k=1}^{K_0} \|\delta_k\|^2 = 1$.

By the permutation, such a parametrization is locally conic and identifiable. And then, the penalized quasi-likelihood function (2.4) can be rewritten as

$$Q_{\mathrm{P}}(\eta,\gamma) \equiv \sum_{i=1}^{n} \log\{\Psi(\eta,\gamma;Y_i \mid X_i)\} - n\lambda \sum_{k=1}^{K} \{\log(\epsilon+\pi_k) - \log(\epsilon)\}. \tag{3.2}$$

To establish the model selection consistency of the proposed method, we need the following additional conditions:

C5 There exists a positive constant ε such that $g(X_{ij}^T\beta_k)$ and $V(g(X_{ij}^T\beta_k))$ are bounded on $\mathcal{B} = \{\beta : \|\beta - \beta_0\| \le \varepsilon\}$ uniformly in $i = 1,\ldots,n$, $j = 1,\ldots,m_i$, $k = 1,\ldots,K$.

C6 Let $\Sigma_{ik} = \mathrm{cov}(Y_i \mid X_i, u_i = k)$, and V_{ik} be a $m_i \times m_i$ diagonal matrix with jth element σ_{ijk}^2. Both the eigenvalues of Σ_{ik} and V_{ik} are uniformly bounded away from 0 and infinity.

Condition C5 is analogous to conditions (A2) and (A6) in Wang (2011), which is generally satisfied for marginal models. For example, when the marginal model follows a Poisson regression, $V(g(X_{ij}^T\beta_k)) = g(X_{ij}^T\beta_k) = \exp(X_{ij}^T\beta_k)$'s are uniformly bounded around on $\mathcal{B}$. Condition C6 is similar to conditions (C3) and (C4) in Huang et. al. (2007), which ensures the non-singularity of the covariance matrices and the working covariance matrices.

THEOREM 2. *Under conditions C1-C6, if* $\lim_{n\to\infty}\sqrt{n}\lambda = a$ *and* $\epsilon = o(n^{-1/2}/\log n)$, *where* a *is a constant, there exists a local maximizer* (η,γ) *of (3.2) such that* $\eta = O_p(n^{-1/2})$, *and for such local maximizer, we have* $\widehat{K} \xrightarrow{P} K_0$.

Theorem 2 indicates that by choosing an appropriate tuning parameter λ and a small constant ϵ, the proposed method (2.4) can select the number of mixing components consistently.

4 Implementation and tuning parameter selection

In this section, we propose an algorithm to implement the proposed method (2.4) and a procedure to select the tuning parameter λ.

4.1 Modified EM Algorithm

Since the membership of each subject is unknown, it is natural to use EM algorithm to implement (2.4). But note that the criterion (2.4) is a function unrelated to different dispersion parameters $\phi_k, k = 1, \ldots, K$, therefore, the naive EM algorithm may decrease the classification accuracy for the observations by ignoring the within-component dispersion. Therefore, we here propose a modified EM algorithm in consideration of different component dispersion.

Let u_{ik} denote the indicator of whether the ith subject is in the kth class. That is, $u_{ik} = 1$ if the ith subject belongs to the kth component, and $u_{ik} = 0$ otherwise. If the missing data $\{u_{ik}, i = 1, \ldots, n, k = 1, \ldots, K\}$ were observed, the penalized quasi-likelihood function for the complete data is given by

$$\sum_{i=1}^{n}\sum_{k=1}^{K} u_{ik}\left\{\log \pi_k + \sum_{j=1}^{m_i} q(g(X_{ij}^T\beta_k); Y_{ij})\right\} - n\lambda\sum_{k=1}^{K}\{\log(\epsilon + \pi_k) - \log(\epsilon)\}. \tag{4.1}$$

Denote $\Theta = (\pi^T, \beta^T, \phi^T)^T$ as the vector of all parameters in the K-component marginal mixture regression model (2.1)-(2.2) with $\beta = (\beta_1^T, \ldots, \beta_K^T)^T$. In the E-step, given the current estimate $\Theta^{(t)} = (\pi^{(t)T}, \beta^{(t)T}, \phi^{(t)T})^T$, we impute values for the unobserved u_{ik} by

$$\widehat{u}_{ik}^{(t+1)} = \frac{\pi_k^{(t)} \exp\left\{\sum_{j=1}^{m_i} \tilde{q}(g(X_{ij}^T\beta_k^{(t)}), \phi_k^{(t)}; Y_{ij})\right\}}{\sum_{l=1}^{K} \pi_l^{(t)} \exp\left\{\sum_{j=1}^{m_l} \tilde{q}(g(X_{ij}^T\beta_l^{(t)}), \phi_l^{(t)}; Y_{ij})\right\}},$$

where $\tilde{q}(\mu, \phi; y) = \int_y^\mu \frac{y-t}{\phi V(t)}dt$. Plugging them into (4.1), we obtain the function

$$\sum_{i=1}^{n}\sum_{k=1}^{K} \widehat{u}_{ik}^{(t+1)}\left\{\log \pi_k + \sum_{j=1}^{m_i} q(g(X_{ij}^T\beta_k); Y_{ij})\right\} - n\lambda\sum_{k=1}^{K}\{\log(\epsilon + \pi_k) - \log(\epsilon)\}. \tag{4.2}$$

In the M-step, the goal is to update $\pi^{(t)}$ and $\beta^{(t)}$ by maximizing (4.2) with the constraint $\sum_{k=1}^{K} \pi_k = 1$ and update $\phi^{(t)}$ by the residual moment method. Specifically, to update $\pi^{(t)}$, we solve the following equations

$$\frac{\partial}{\partial \pi}\left\{\sum_{i=1}^{n}\sum_{k=1}^{K} \widehat{u}_{ik}^{(t+1)} \log \pi_k - n\lambda\sum_{k=1}^{K}\{\log(\epsilon + \pi_k) - \log(\epsilon)\} - \xi(\sum_{k=1}^{K}\pi_k - 1)\right\} = 0,$$

where ξ is the Lagrange multiplier. Then, when ϵ is very close to zero, it gives

$$\pi_k^{(t+1)} = \max\left\{0, \frac{1}{1 - \lambda K}\left[\frac{1}{n}\sum_{i=1}^{n} \widehat{u}_{ik}^{(t+1)} - \lambda\right]\right\}, k = 1, \ldots, K. \tag{4.3}$$

$\beta_k^{(t)}$ can be updated by solving the following equations

$$\sum_{i=1}^{n}\sum_{j=1}^{m_i}\widehat{u}_{ik}^{(t+1)}g'(X_{ij}^T\beta_k)X_{ij}\frac{Y_{ij}-g(X_{ij}^T\beta_k)}{V(g(X_{ij}^T\beta_k))}=0,$$

where $g'(\cdot)$ is the first derivative of g, $k=1,\ldots,K$. And using the residual moment method, we update $\phi^{(t)}$ as follows

$$\phi_k^{(t+1)}=\sum_{i=1}^{n}\frac{\widehat{u}_{ik}^{(t+1)}}{\sum_{i'=1}^{n}m_{i'}\widehat{u}_{i'k}^{(t+1)}}\sum_{j=1}^{m_i}\frac{\{Y_{ij}-g(X_{ij}^T\beta_k^{(t)})\}^2}{V(g(X_{ij}^T\beta_k^{(t)}))},k=1,\ldots,K.$$

Remark 3. *In the initial step, we pre-specify a large number of components, and once a mixing proportion is shrunk to zero by (4.3), the corresponding parameters in this component are set to zero and fewer components are kept for the remaining EM iterations. Here we use the same notation K for the whole process. In practice, during the iterations, K becomes smaller and smaller until the algorithm converges.*

Remark 4. *Although in theory we require $\epsilon=o(n^{-1/2}/\log n)$, we can update π using (4.3) without choosing ϵ in practice.*

4.2 Turning Parameter Selection and Classification Rule

In terms of selecting the tuning parameter λ, we follow the suggestion in Wang, Li, and Tsai (2007) and use a BIC-type criterion:

$$\mathrm{BIC}(\lambda)=-2\sum_{i=1}^{n}\log\left[\sum_{k=1}^{\widehat{K}}\widehat{\pi}_k\exp\left\{\sum_{j=1}^{m_i}\tilde{q}(g(X_{ij}^T\widehat{\beta}_k),\widehat{\phi}_k;Y_{ij})\right\}\right]+\widehat{K}(p+2)\log n, \quad (4.4)$$

where $\widehat{K}$ and $\widehat{\beta}$ are estimators of K_0 and β_0 by maximizing (2.4) for a given λ.

Let $\widehat{K}$, $\widehat{\pi}$, $\widehat{\beta}$, and $\widehat{\phi}$ be the final estimators of the number of components, the mixing proportions and unknown parameters, respectively. Then, in the sense of clustering, a subject can be assigned to the class whose empirical posterior is the largest. For example, a subject (Y^*,X^*) with m times observations is assigned to the class

$$k^*=\arg\max_{1\le k\le\widehat{K}}\widehat{\pi}_k\exp\left\{\sum_{j=1}^{m}\tilde{q}(g(X_{ij}^{*T}\widehat{\beta}_k),\widehat{\phi}_k;Y_{ij}^*)\right\}. \quad (4.5)$$

Consequently, a nature predictor of Y^* is given by $g(X^T\widehat{\beta}_{k^*})$.

REMARK 5. *One may claim that $\widehat{\beta}$ would loss some efficiency if the within-subject correlation is strong. It would be better to incorporate correlation information to gain estimation efficiency. However, a correlation analysis would lead to additional computational cost and increase the chance of the convergence problem for the proposed modified EM algorithm. In practice, we suggest to estimate β once again given the component information derived from (2.4). Specifically, we first fit the mixture regression model (2.1)-(2.2) and cluster samples into $\widehat{K}$ classes by (4.5); then, in each class, the marginal generalized linear model is estimated by applying GEE with a working correlation structure. It is expected that this two-step technique may improve the estimation efficiency if the correlation of the longitudinal data is strong and the working structure is correctly specified.*

5 Simulation studies

In this section, we conduct a set of Monte Carlo simulation studies to assess the finite sample performance of the proposed method. The maximum initial number of clusters is set to be ten, the initial value for the modified EM algorithm is estimated by K-means clustering and the tuning parameter λ is selected by the proposed BIC criterion (4.4). To test the classification accuracy and estimation accuracy, we conduct 200 replications and compare the method (2.4) with the two-step method mentioned in Remark 5 and the QIFC method proposed by Wang and Qu (2014). QIFC is a supervised classification technique for longitudinal data. To permit comparison, we assume that the true number of components, the true class label and the true within-subject correlation are known for the QIFC method. We denote the proposed method and the two-step method as PQL and PQL2 in the following, respectively.

Example 1. We generate correlated count outcomes from a two-component overdispersed Poisson mixture with mixing proportions $\pi_1 = 1/3$ and $\pi_2 = 2/3$. For component 1, the mean function of repeated measurements Y_{ij} is

$$\log(\mu_{ij1}) = 3X_{ij1} - X_{ij2} + X_{ij3}, \qquad i = 1, \ldots, n_1, j = 1, \ldots, m_i,$$

and the marginal variance is $\phi_1\mu_{ij1} = \text{var}(Y_{ij}) = 2\mu_{ij1}$. The correlation structure within a subject is AR(1) with correlation coefficient ρ. For component 2, Y_i has the same correlation matrix as in component 1, except that

$$\log(\mu_{ij2}) = 4 - 2X_{ij1} + X_{ij3}, \qquad i = 1, \ldots, n_2, j = 1, \ldots, m_i,$$

with dispersion parameter $\phi_2 = 1$. The number of repeated measurements m_i is randomly drawn from a Poisson distribution with mean 3 and increased by 2, and the sample size is $n = 150$. Covariates X_{ijp} are generated independently from uniform distribution $U(0, 1)$. Two values of ρ are considered, $\rho = 0.3$ and 0.6, to represent different correlation magnitude.

To measure the performance of the proposed tuning parameter selector (4.4), we show the histograms of the estimated component numbers and report the percentage of selecting the correct number of components. With respect to classification, we generate 100 new subjects from each component with the same setting as in each configuration and measure the performance in terms of the misclassification error rate. We summarize the median and the 95% confidence interval of misclassification error rate from a model with correctly identified K_0 for PQL and PQL2 and report these quantities for QIFC as well. To measure the performance of the proposed estimators, the median and the mean squared error (MSE) of the estimators are reported provided that the number of components K_0 is correctly identified. Correspondingly, the median and the MSE of the estimated QIFC estimators are also summarized as a benchmark for comparison. Note that the label switching might arise in practice. Yao (2015) and Zhu and Fan (2016) proposed many feasible labeling methods and algorithms. In our simulation studies, we solve the label switching by putting an order constraint on components' mean parameters.

Figure 1(a) and (b) depict the histograms of the estimated component numbers with different correlation magnitude. It shows that the proposed PQL method with the BIC tuning parameter selector can identify the correct number of components at least with probability 0.945, which is in accordance with the model selection consistency in Theo-

rem 2.

When the number of components is correctly identified, Table 4(a) reports the median and the 95% confidence interval of the misclassification error rate from the model-based clustering. We can see that the proposed methods perform better than QIFC with relatively smaller misclassification error rate. Since QIFC is proved asymptotically optimal in terms of misclassification error rate (see Theorem 1 in Wang and Qu, 2014), the observations in Table 4(a) imply the optimality of the proposed methods in terms of misclassification error rate numerically. Further, in terms of parameter estimation, Table 5 summarizes the estimation of the unknown parameters for a two-component overdispersed Poisson mixture model under different correlation magnitude. We observe that the PQL estimators are consistent as discussed in Theorem 1. It is not surprising that, for regression parameters, PQL performs not as well as QIFC with larger MSE especially when the within-subject correlation is high, since QIFC estimators are oracle by assuming the known class memberships and the true within-subject correlation structure. It implies that ignoring the working correlation would affect the efficiency of parameter estimation. However, we can improve the estimating efficiency if the correct correlation information is incorporated. This is reflected in that the PQL2 estimators have much smaller MSEs compared with the PQL estimators. Indeed, the PQL2 performs similarly to the QIFC estimator.

In addition, combining Table 4(a) and Table 5, we can observe that the two-step technique is able to improve the estimation efficiency for the mean regression parameters without reducing the classification accuracy, which validate our guess in Remark 5 numerically. In general, when the within-subject correlation is large, it is recommended to use PQL2 to provide more predictive power by utilizing the within-subject correlation information.

Example 2. In this example, we consider a mixture of three linear regression models. We

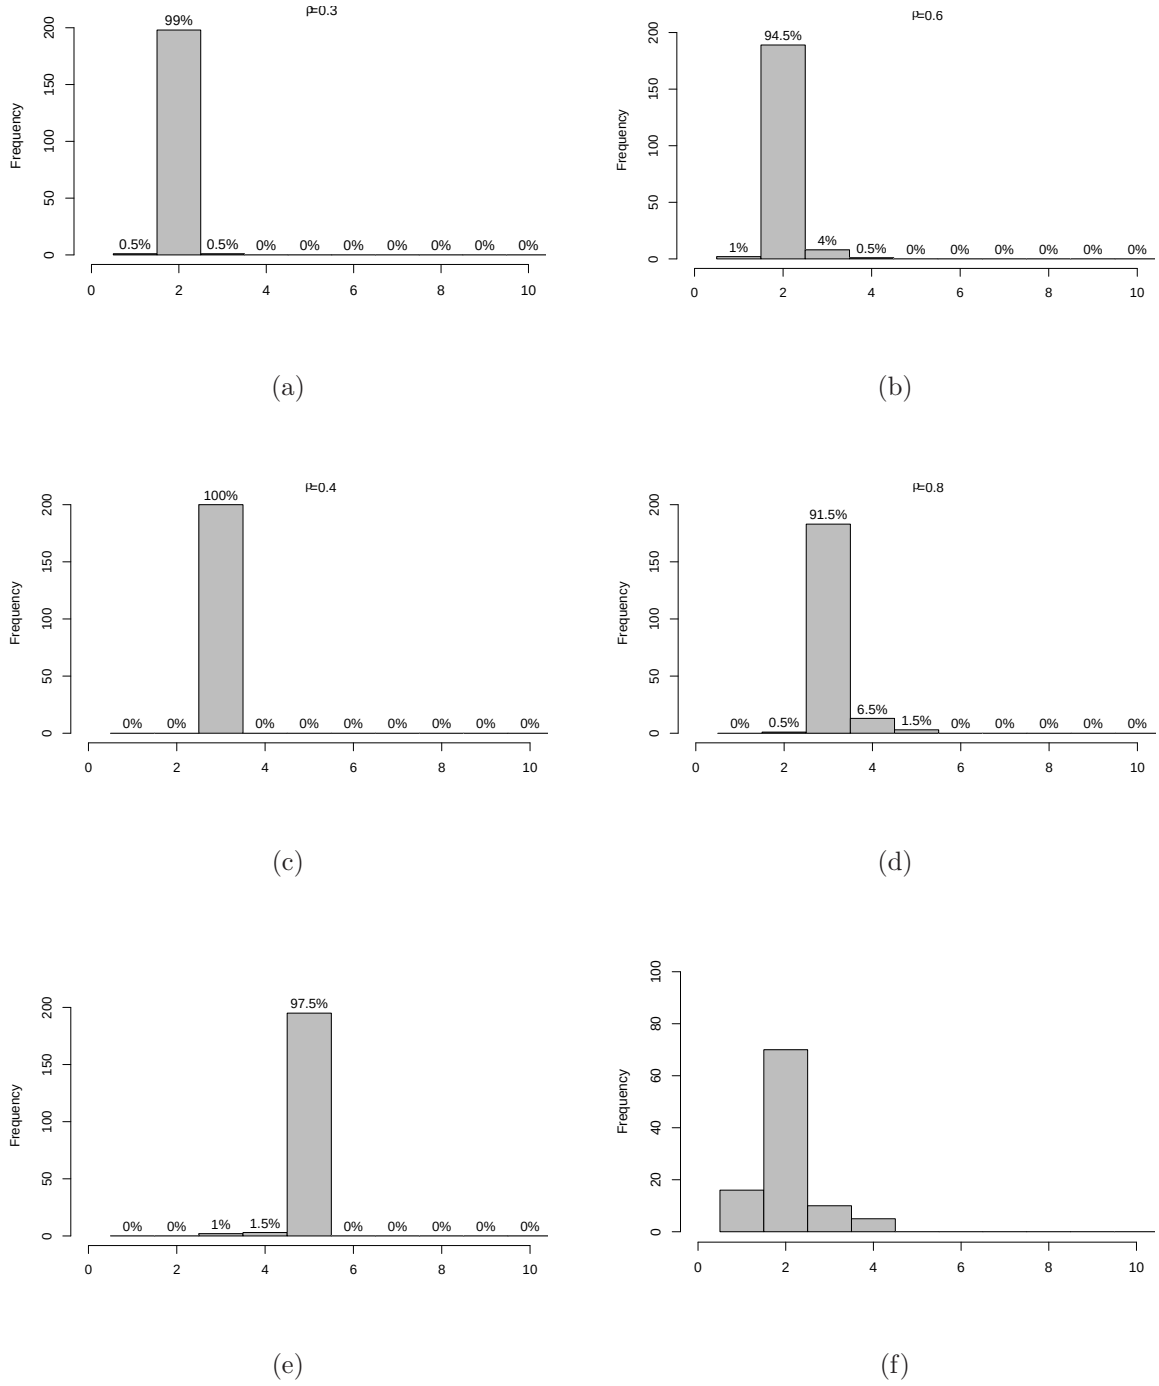

Figure 1: (a) and (b) are histograms of estimated numbers of components by the proposed PQL method in Example 1 with different correlation magnitude. (c) and (d) are histograms of estimated numbers of components by the proposed PQL method in Example 2 with different correlation magnitude. (e) is the histogram of estimated numbers of components by the proposed PQL method in Example 3. The value on the top of each bar is the percentage of selecting the corresponding number of components. (f) is the histogram of estimated number of components based on 100 replications in PBC data.

Table 1: The median and the 95% confidence interval (CI) of total misclassification error rate in simulation studies. The values of median in Examples 1 and 2 are times 100. For the proposed PQL and PQL2 methods, the results below are summarized based on the models with correctly specified K_0 in 200 replications.

(a)Example 1		Criterion	PQL	PQL2	QIFC
	$\rho = 0.3$	median	0.000	0.000	0.500
		CI	(0.000, 0.010)	(0.000, 0.010)	(0.000, 0.025)
	$\rho = 0.6$	median	0.000	0.000	0.700
		CI	(0.000, 0.010)	(0.000, 0.010)	(0.000, 0.020)
(b)Example 2		Criterion	PQL	PQL2	QIFC
	$\rho = 0.4$	median	0.333	0.333	0.343
		CI	(0.000, 0.010)	(0.000, 0.010)	(0.000, 0.013)
	$\rho = 0.8$	median	0.333	0.333	0.352
		CI	(0.000, 0.010)	(0.000, 0.010)	(0.000, 0.013)
(c)Example 3		Criterion	PQL	PQL2	QIFC
		median	0.208	0.208	0.264
		CI	(0.200, 0.218)	(0.200, 0.218)	(0.214, 0.304)

Table 2: The median and the mean squared error (MSE) on parameter estimation based on 200 replicates in Example 1. The values of MSE are times 100. For the proposed PQL and PQL2 methods, the results below are summarized based on the models with correctly specified K_0 in 200 replications.

	$\rho = 0.3$						$\rho = 0.6$					
	PQL		PQL2		QIFC		PQL		PQL2		QIFC	
	median	MSE	median	MSE	median	MSE	median	MSE	median	MSE	median	MSE
β_{10}	-0.005	1.082	0.001	0.987	-0.012	1.061	0.022	1.413	0.018	1.042	-0.015	1.088
β_{11}	2.998	1.229	2.994	1.109	3.011	1.247	3.005	1.238	2.990	0.753	3.014	0.911
β_{12}	-0.993	0.738	-0.992	0.657	-0.994	0.701	-1.007	0.825	-1.015	0.462	-1.022	0.524
β_{13}	0.988	0.679	0.989	0.696	0.989	0.733	0.990	0.777	0.990	0.457	0.999	0.431
β_{20}	3.998	0.119	3.998	0.104	3.994	0.110	4.001	0.101	3.997	0.063	3.997	0.071
β_{21}	-1.994	0.144	-1.996	0.119	-1.997	0.139	-1.997	0.148	-1.997	0.082	-1.998	0.105
β_{22}	0.002	0.092	0.001	0.088	0.001	0.094	0.004	0.119	0.001	0.066	0.001	0.075
β_{23}	1.002	0.149	1.003	0.128	1.004	0.136	0.998	0.131	0.999	0.071	0.999	0.089
ϕ_1	1.974	2.713	1.993	2.824	2.003	2.868	1.952	4.891	1.972	5.259	1.986	5.448
ϕ_2	0.980	1.167	0.996	1.174	0.986	1.158	0.956	1.577	0.982	1.540	0.969	1.520
π_1	0.332	0.011	0.328	0.011	–	–	0.314	0.030	0.314	0.030	–	–
π_2	0.668	0.011	0.668	0.011	–	–	0.686	0.030	0.686	0.030	–	–

generate 300 samples with mixing proportions $\pi_1 = \pi_2 = 0.3$ and $\pi_3 = 0.4$. Conditional on the class label u_i, we generate the response vector Y_i from three multivariate normal distributions:

$$Y_i \mid u_i = k \sim MVN(\beta_{k0} + X_i\beta_k, \sigma_k^2 R_i^{(k)}(\rho)), k = 1, 2, 3$$

where the component's regression parameters are set as $(\beta_{10}, \beta_1^T)^T = (2, 1, -1, 1.5, 1)$, $(\beta_{20}, \beta_2^T)^T = (-4, 2, 1, -2, 0)$ and $(\beta_{30}, \beta_3^T)^T = (-2, -2, 1, 0, 1)$, the variance parameters are set as $(\sigma_1^2, \sigma_2^2, \sigma_3^2) = (0.5, 0.3, 0.1)$. The number of repeated measurements m_i are drawn as in Example 1. And similar to Example 1, we consider data from a single correlation structure, AR(1). That is, $R_i^{(1)}(\rho) = R_i^{(2)}(\rho) = R_i^{(3)}(\rho) = R_{i\mathrm{AR}(1)}(\rho)$. Two values $\rho = 0.4$ and 0.8 are considered to represent moderate and high correlation. To mimic the real data, the correlated covariates $(X_{ij1}, X_{ij2}, X_{ij3})^T$ are generated independently from the trivariate normal distribution with mean 0.5, marginal variance 0.2 and AR(1) corre-

lation structure with auto-correlation coefficient 0.6, and the discrete covariates X_{ij4} are generated independently from the Bernoulli distribution $B(1, 0.5)$.

Figure 1(c) and (d) show that our proposed method can identify the correct model in more than 90% cases even for the highly correlated continuous data ($\rho = 0.8$). And once the model is correctly selected, the classification accuracy is quite satisfactory. Table 1(b) implies that our approaches provide more predictive power, especially for large within-subject correlation. Table 1 in Appendix B indicates that our modified EM algorithm gives consistent estimates for parameters and mixing proportions by considering the within-class dispersions. Similar to that in Example 1, when the within-subject correlation is large, the PQL2 approach enhances the estimation efficiency by incorporating the correlations within each subject while retaining the class membership prediction accuracy.

Example 3. In the third example, we consider a five-component Gaussian mixture of AR(1), exchangeable (CS), and independence (IND). This is a more challenging example with more components but having different correlation structures. Specifically, we generate 500 samples with mixing proportions $\pi_1 = \pi_2 = 0.25$, $\pi_3 = \pi_4 = 0.15$ and $\pi_5 = 0.2$. Conditional on the class label u_i, the response vector Y_i is generated from five multivariate normal distributions:

$$Y_i \mid u_i = k \sim MVN(\beta_{k0} + X_i\beta_k, \sigma_k^2 R_i^{(k)}), k = 1, \cdots, 5$$

where the component's regression parameters are set as $(\beta_{10}, \beta_1^T)^T = (2, 1, -1, 1.5, 1)$, $(\beta_{20}, \beta_2^T)^T = (-4, 2, 1, -2, 0)$, $(\beta_{30}, \beta_3^T)^T = (-2, -2, 1, 0, 1)$, $(\beta_{40}, \beta_4^T)^T = (0, 1, 0, 1, 1)$, and $(\beta_{50}, \beta_5^T)^T = (-4, 0, -1, -1, -1.5)$, the correlation structures are set as $R_i^{(1)} = R_{iAR(1)}(0.6)$, $R_i^{(2)} = R_{iAR(1)}(0.6)$, $R_i^{(3)} = R_{iCS}(0.3)$, $R_i^{(4)} = R_{iCS}(0.3)$, $R_i^{(5)} = R_{iIND}$, and the variance parameters are set as $(\sigma_1^2, \sigma_2^2, \sigma_3^2, \sigma_4^2, \sigma_5^2) = (0.5, 0.3, 0.1, 0.15, 0.6)$. The number of repeated measurements m_i and the covariates are generated as in Example 2.

Figure 1(e) draws the histogram of estimated numbers of components and Table 1(c) summarizes the classification results in terms of misclassification error rate. Though PQL uses a single correlation structure (IND), it is able to identify the correct number

of components with high probability, and give more accurate prediction of the class's membership compared with QIFC. Table 2 in Appendix B indicates, across different finite mixture correlation models, PQL estimators are still consistent. It may loss some efficiency, but can be improved by PQL2.

6 Application to primary biliary cirrhosis data

In this section, we apply the proposed method to study a doubled-blinded randomized trail in primary biliary cirrhosis (PBC) conducted by the Mayo Clinic between 1974 and 1984 (Dickson, Grambsch, Fleming, Fisher, and Langworthy, 1989).

This data set consists of 312 patients who consented to participate in the randomized placebo controlled trial of the drug D-penicillamine. Each patient was supposed to have measurements taken at 6 months, 1 year, and annually thereafter. However, a sizable portion of patients missed their measurements because of worsening medical condition of some labs, which resulted in an unbalanced data structure. The serum bilirunbin level, date of visit, status and other histologic parameters were recorded for each patient on each day of visit. PBC is a rare but fatal chronic cholestatic liver disease, with a prevalence of about 50-cases-per-millon-population. Affected patients are typically middle-aged women. As in this data set, the sex ratio is $9.3 : 1$ (women to men), where the median age of women patients is 49 years old. Identification of PBC is crucial to balancing the need for medical treatment to halt disease progression and extend survival free of liver transplantation, while minimizing drug-induced toxicities. Therefore, one goal of the study was to investigate the relationship between marker serum bilirunbin level and other covariates of interest.

We set the log-transformed serum bilirunbin level as the response variable, since the original level has positive observed values. Figure 2(a) depicts the plot of a set of observed transformed longitudinal profiles of serum bilirunbin marker. It shows that the variability may be large for different patients. In addition, two sample t-test indicates that there exists significant difference in means of age between placebo and D-penicillamine group

(p-value = 0.02). Therefore, we consider the marginal semiparametric mixture regression model (2.1)-(2.2) with the identity link. The mean structure in the kth component takes the form

$$\mathrm{E}(Y_{ij}) = \beta_{1k}\mathrm{Trt}_{ij} + \beta_{2k}\mathrm{Age}_{ij} + \beta_{3k}\mathrm{Gender}_{ij} + \beta_{4k}\mathrm{Time}_{ij} + \beta_{5k}\mathrm{Trt}_{ij} * \mathrm{Time}_{ij}, \qquad (6.1)$$

and the marginal variance is assumed as $\mathrm{var}(Y_{ij}) = \sigma_k^2$, $i = 1, \ldots, 312$, $j = 1, \ldots, m_i$, $k = 1, \ldots, K$, where variable Trt is a binary variable with 0 for D-penicillamine and 1 for placebo, variable Gender is binary with 0 for male and 1 for female, and variable Time is the number of months between enrollment and this visit date.

We first standardize data so that there is no intercept term in model (6.1). Then, we apply the proposed method to simultaneously select the number of components and to estimate the mixing proportions and unknown parameters. As in the simulation studies, the maximum initial number of clusters is set to be ten and the initial value for the modified EM algorithm is estimated by K-means clustering. The estimated number of components is $\widehat{K} = 2$. To check its stability, we run our proposed method 100 replications. For each run, we randomly draw 80% of patients for each status. Figure 1(f) shows that our proposed method selects two components with high probability. The resulting estimators of parameters and mixing proportions along with the 95% confidence intervals are shown in Table 3.

Further, classification of patients according to criterion (4.5) leads to 172 patients being classified in class 1 and 140 patients in class 2. According to the estimates of parameters in Table 3, it implies that class 1 is characterized by a relatively lower baseline bilirubin level. Moreover, note that $\widehat{\beta}_{14} = -0.216 < -\widehat{\beta}_{15} = -0.115$ and $\widehat{\beta}_{24} = 4.047 > -\widehat{\beta}_{25} = 3.929$, which implies that the bilirubin level increases more slowly over time in class 1. Therefore, from the clinical point of view, class 1 should correspond to patients with a better prognosis compared to class 2. To confirm this conclusion, Kaplan-Meier estimates of the survival probabilities are calculated based on data from patients classified in each class. We can see that, from Figure 2(b), the survival prognosis of class 1 is indeed much

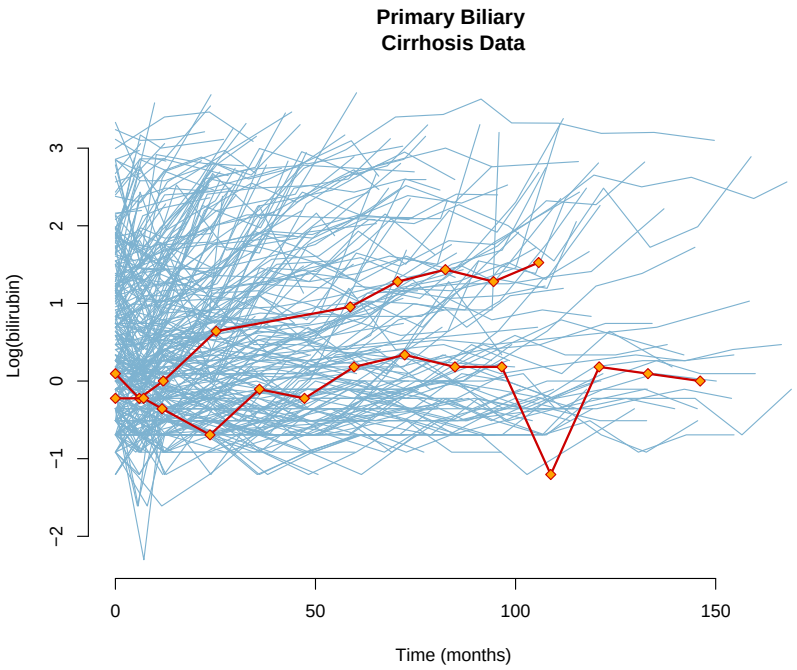

(a)

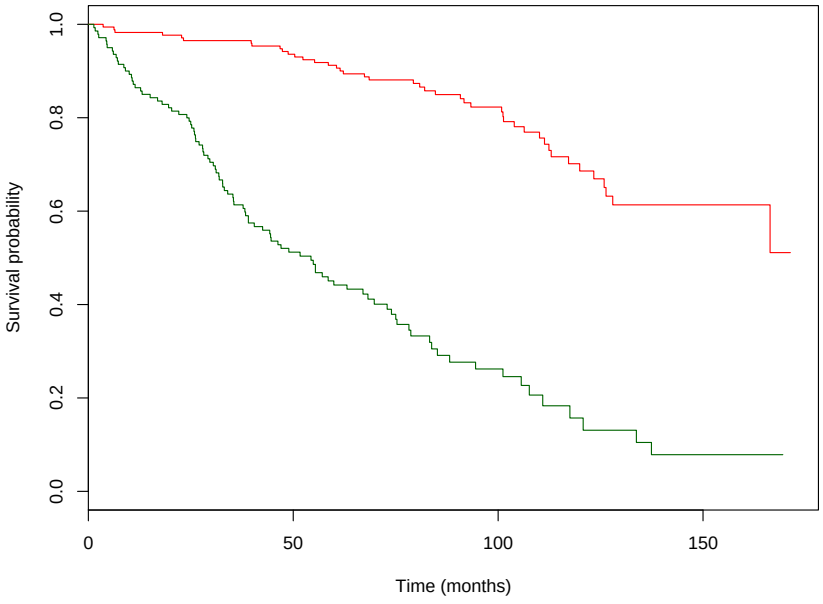

(b)

Figure 2: (a): Observed transformed longitudinal profiles of serum bilirunbin marker. The red lines are profiles of two selected patients (id 2 and 34). (b): The Kaplan-Meier estimate of survival curves for two classes (class 1: red, class 2: green).

Table 3: Parameter estimation with $\widehat{K} = 2$ along with the 95% confidence intervals.

Parameters	Component 1		Component 2	
	Estimation	Confidence interval (95%)	Estimation	Confidence interval (95%)
mixing proportions	0.554	(0.442, 0.641)	0.446	(0.359, 0.558)
Trt	0.252	(0.142, 0.339)	0.955	(0.738, 1.074)
Age	0.057	(-0.010, 0.135)	0.167	(0.066, 0.265)
Gender	0.061	(0.016, 0.193)	0.143	(0.066, 0.234)
Time	-0.216	(-0.288, -0.105)	4.047	(3.098, 4.507)
Trt$*$Time	0.115	(0.016, 0.177)	-3.929	(-4.479 -2.959)
σ^2	0.609	(0.467, 0.754)	1.140	(0.833, 1.531)

better than that of class 2 with the estimated 5-year survival probability in class 1 of 0.912 compared to 0.442 in class 2, and the 10-year survival probabilities 0.686 and 0.157 in classes 1 and 2, respectively.

7 Conclusion

In this paper, we have proposed a penalized method for learning mixture regression models from longitudinal data which is able to select the number of components in an unsupervised way. The proposed method only requires the first two moment conditions of the model distribution, and thus is suitable for both the continuous and discrete responses. It penalizes the logarithm of mixing proportions, which allows one to simultaneously select the number of components and to estimate the mixing proportions and unknown parameters. Theoretically, we have shown that our proposed approach can select the number of components consistently for general marginal semiparametric mixture regression models. And given the number of components, the estimators of mixing proportions and regression parameters are root-n consistent and asymptotically normal.

To improve the classification accuracy, a modified EM algorithm has been proposed by considering the within-component dispersion. Simulation results and the real data

analysis have shown its convergence, but further theoretical investigation is needed. And we have introduced a BIC-type method to select the tuning parameter automatically. Numerical studies show it works well, while the theoretical consistency deserves a further study.

Another issue is the consideration of the within-subject correlation. The proposed penalized approach is introduced under the working independence correlation. Simulation results have implied that it may loss some estimation efficiency, especially when the within-subject correlation is large. Therefore, we suggest a two-step technique to refine the estimates. Simulations show that the efficiency improvement is significant if the correlation information is incorporated and the working structure is correctly specified. It would be worthwhile to systematically study the unsupervised learning of mixtures by incorporating correlations.

References

[1] Booth, J.G., Casella, G., and Hobert, J.P. (2008). Clustering using objective functions and stochastic search. *Journal of the Royal Statistical Society* **B70,** 119–139.

[2] Celeux, G., Martin, O., and Lavergne, C. (2005). Mixture of linear mixed models for clustering gene expression profiles from repeated microarray experiments. *Statistical Modelling* **5,** 243–267.

[3] Chen, J. and Khalili, A. (2008). Order selection in finite mixture models with a non-smooth penalty. *Journal of the American Statistical Association* **104,** 187–196.

[4] Dacunha-Castelle, D. and Gassiat, K. (1997). Testing in locally conic models and application to mixture models. *ESAIM: Probability and Statistics* **1,** 285–317.

[5] Dacunha-Castelle, D. and Gassiat, K. (1999). Testing the order of a model using locally conic parametrization: population mixtures and stationary ARMA processes. *The Annals of Statistics* **27,** 1178–1209.

[6] Dasgupta, A. and Raftery, A.E. (1998). Detecting features in spatial point processes with clutter via model-based clustering. *Journal of the American Statistical Association* **93,** 294–302.

[7] De la Cruz-Mesía, R., Quintana, F.A., and Marshall, G. (2008). Model-based clustering for longitudinal data. *Computational Statistics & Data Analysis* **52,** 1441–1457.

[8] Dickson, E.R., Grambsch, P.M., Fleming, T.R., Fisher, L.D., and Langworthy, A. (1989). Prognosis in primary biliary cirrhosis: Model for decision making. *Hepatology* **10,** 1–7.

[9] Erosheva, E.A., Matsueda, R.L., and Telesca, D. (2014). Breadking bad: two decades of life-course data analysis in criminology, developmental psychology, and beyond. *Annual Review of Statistics and Its Application* **1,** 301–332.

[10] Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association* **96,** 1348–1360.

[11] Fraley, C. and Raftery, A.E. (2002). Model-based clustering discriminant analysis and density estimation. *Journal of the American Statistical Association* **97,** 611–631.

[12] Genolini, C. and Falissard, B. (2010). KmL: k-means for longitudinal data. *Computational Statistics* **25,** 317–328.

[13] Heinzl, F. and Tutz, G. (2013). Clustering in linear mixed models with approximate Dirichlet process mixtures using EM algorithm. *Statistical Modelling* **13,** 41–67.

[14] Heinzl, F. and Tutz, G. (2014). Clustering in linear-mixed models with a group fused lasso penalty. *Biometrical Journal* **56,** 44–68.

[15] Hennig, C. (2004). Breakdown points for maximum likelihood estimators of location-scale mixtures. *The Annals of Statistics* **32,** 1313–1340.

[16] Huang, J.Z., Zhang, L., and Zhou, L. (2007). Efficient estimation in marginal partially linear models for longitudinal/clustered data using splines. *Scandinavian Journal of Statistics* **34,** 451–477.

[17] Huang, T., Peng, H., and Zhang, K. (2016). Model selection for Gaussian mixture models. *Statistica Sinica*, in press.

[18] Keribin, C. (2000). Consistent estimation of the order of mixture models. *Sankhyā* **62,** 49–66.

[19] Komárek, A. and Komárková, L. (2013). Clustering for multivariate continuous and discrete longitudinal data. *The Annals of Applied Statistics* **7,** 177–200.

[20] Komárek, A. and Lesaffre, E. (2008). Generalized linear mixed model with a penalized Gaussian mixture as a random effects distribution. *Computational Statistics & Data Analysis* **52,** 3441–3458.

[21] Leroux, B. (1992). Consistent estimation of a mixing distribution. *The Annals of Statistics* **20,** 1350–1360.

[22] Liang, K.Y. and Zeger, S.L. (1986). Longitudinal data analysis using generalised linear models. *Biometrika* **73,** 12–22.

[23] Maruotti, A. (2011). Mixed hidden Markov models for longitudinal data: an overview. *International Statistical Review* **79,** 427–454.

[24] McNicholas, P.D. and Murphy, T.B. (2010). Model-based clustering of longitudinal data. *The Canadian Journal of Statistics* **38,** 153–168.

[25] Pickles, A. and Croudace, T. (2010). Latent mixture models for multivariate and longitudinal outcomes. *Statistical Methods in Medical Reseaerch* **19,** 271–289.

[26] Roeder, K. and Wasserman, L. (1997). Practical density estimation using mixtures of normals. *Journal of the American Statistical Association* **92,** 894–902.

[27] Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics* **6,** 461–464.

[28] Wang, H., Li, R., and Tsai, C.L. (2007). Tuning parameter selelctors for the smoothly clipped absolute deviation method. *Biometrika* **94,** 553–568.

[29] Wang, L. (2011). GEE analysis of clustered bianry data with diverging number of covariates. *The Annals of Statistics* **39,** 389–417.

[30] Wang, X. and Qu, A. (2014). Efficient classification for longitudinal data. *Computatioinal Statistics and Data Analysis* **78,** 119–134.

[31] Xu, P., Zhang, J., Huang, X., and Wang, T. (2016). Efficient estimation of marginal generalized partially linear single-index models with longitudinal data. *TEST* **25,** 413–431.

[32] Xu, P. and Zhu, L. (2012). Estimation for a marginal generalized single-index longitudinal model. *Journal of Multivariate Analysis* **105,** 285–299.

[33] Yao, W. (2015). Label switching and its solutions for frequentist mixture model. *Journal of Statistical Computation and Simulation* **85,** 1000–1012.

[34] Zhu, W. and Fan, Y. (2016). Relabelling algorithms for mixture models with applications for large data sets. *Journal of Statistical Computation and Simulation* **86,** 394–413.

[35] Hastie, T., Tibshirani, R., and Friedman, J. (2001). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York: Springer.

8 Appendix

8.1 A. Proofs of Theorems

Proof of Theorem 1

Recall $\Psi(\theta; Y_i|X_i) = \sum_{k=1}^{K} \pi_k \exp\left\{\sum_{j=1}^{m_i} q(g(X_{ij}^T\beta_k); Y_{ij})\right\}$, and $\psi(\theta; Y_i \mid X_i) = \log(\Psi(\theta; Y_i \mid X_i))$. Let $\widehat{\theta}^* = \sqrt{n}(\widehat{\theta} - \theta_0)$. Then, $\widehat{\theta}^*$ maximizes

$$Q_n(\widehat{\theta}^*) = \sum_{i=1}^{n} \{\psi(n^{-1/2}\widehat{\theta}^* + \theta_0; Y_i \mid X_i) - \psi(\theta_0; Y_i \mid X_i)\}.$$

An application of Taylor expansion yields that

$$\begin{aligned} Q_n(\widehat{\theta}^*) &= \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \frac{\partial \psi(\theta_0; Y_i \mid X_i)}{\partial \theta} \widehat{\theta}^* + \frac{1}{2} \widehat{\theta}^{*T} \left\{ \frac{1}{n} \sum_{i=1}^{n} \frac{\partial^2 \psi(\theta_0; Y_i \mid X_i)}{\partial \theta \partial \theta^T} \right\} \widehat{\theta}^* + o_p(\mathbf{1}) \\ &\equiv D_n \widehat{\theta}^* + \frac{1}{2} \widehat{\theta}^{*T} B_n \widehat{\theta}^* + o_p(\mathbf{1}), \end{aligned} \tag{8.1}$$

where $\mathbf{1}$ is a $pK + (K-1)$ dimensional all-ones vector. It can be shown that $B_n \xrightarrow{P} -B$. Then, by (8.1) and quadratic approximation lemma, we have $\widehat{\theta}^* = B^{-1}D_n + o_p(\mathbf{1})$. Note that $\text{var}(D_n) = A$. And under the regularity conditions, we have $D_n \xrightarrow{L} N(0, A)$. Hence,

$$\widehat{\theta}^* \xrightarrow{L} N(0, B^{-1}AB^{-1}).$$

In order to establish Theorem 2, we need the following lemma first, which can be derived using similar arguments as the proof of Proposition A.1 of Huang et al. (2016).

For a data pair (Y, X) with m times observations, define $\Psi_0(Y \mid X) = \sum_{k=1}^{K_0} \pi_{0k} f(\beta_{0k}; Y \mid X)$. Let $\mathcal{D}$ be the subset of functions of form

$$d(\gamma; Y \mid X) = \sum_{k=1}^{K_0} \pi_{0k} \sum_{j=1}^{p} \delta_{kj} \frac{D_j^1 f(\beta_{0k}; Y \mid X)}{\Psi_0(Y \mid X)} + \sum_{l=1}^{K-K_0} \lambda_l \frac{f(\beta_l; Y \mid X)}{\Psi_0(Y \mid X)} + \sum_{k=1}^{K_0} \rho_k \frac{f(\beta_{0k}; Y \mid X)}{\Psi_0(Y \mid X)},$$

where $D_j^1 f(\beta_{0k}; Y \mid X)$ is the first derivative of $f(\beta_{0k}; Y \mid X)$ for the jth component of β_{0k}.

LEMMA 8.1. *Under conditions C1-C6, $\mathcal{D}$ is a Donsker class.*

Proof: Under conditions C1-C6, it is straightforward to show that $\mathcal{D}$ satisfies conditions P0 and P1 in Dacunha-Castelle and Gassiat (1999) as in Keribin (2000). Then, there exists

a $\Psi_0\nu$-square integrable envelope function $\bar{d}(\cdot)$ such that $|d(\gamma; Y|X)| \leq \bar{d}(Y|X)$. On the other hand, the sequences of coefficients of $d(\gamma; Y|X)$ are bounded under the restrictions imposed on γ. Hence, similar to the proof of Proposition 3.1 in Dacunha-Castelle and Gassiat (1999), we can show that $\mathcal{D}$ has the Donsker property with the bracketing number $N(\varepsilon) = \varepsilon^{-pK}$.

Proof of Theorem 2

In the spirit of proof of Theorem 3.2 in Huang et al. (2016), we divide our proof into two parts. First, we show that there exists a maximizer (η, γ) such that $\eta = O_p(n^{-1/2})$ when $\lambda = a/\sqrt{n}$. It is sufficient to show that, for a large constant C, $Q_{\mathrm{P}}(\eta, \gamma) < Q_{\mathrm{P}}(0, \gamma)$ when $\eta = C/\sqrt{n}$. Let $\eta = C/\sqrt{n}$, and note that

$$\begin{aligned} Q_{\mathrm{P}}(\eta, \gamma) - Q_{\mathrm{P}}(0, \gamma) &= \sum_{i=1}^{n} \{\log \Psi(\eta, \gamma; Y_i \mid X_i) - \log \Psi_0(Y_i \mid X_i)\} \\ &\quad -n\lambda \sum_{l=K-K_0+1}^{K} \{\log(\epsilon + \pi_l) - \log(\epsilon + \pi_{0(l-K+K_0)})\} - n\lambda \sum_{k=1}^{K-K_0} \{\log(\epsilon + \pi_k) - \log(\epsilon)\}. \end{aligned}$$

Then,
$$\begin{aligned} Q_{\mathrm{P}}(\eta, \gamma) - Q_{\mathrm{P}}(0, \gamma) &\leq \sum_{i=1}^{n} \{\log \Psi(\eta, \gamma; Y_i \mid X_i) - \log \Psi_0(Y_i \mid X_i)\} \\ &\quad -n\lambda \sum_{l=K-K_0+1}^{K} \{\log(\epsilon + \pi_l) - \log(\epsilon + \pi_{0(l-K+K_0)})\} := S_1 + S_2. \end{aligned}$$

For S_1, an application of Taylor expansion yields

$$\begin{aligned} S_1 \;=\; & \sum_{i=1}^{n} \frac{\Psi(\eta, \gamma; Y_i \mid X_i) - \Psi_0(Y_i \mid X_i)}{\Psi_0(Y_i \mid X_i)} - \frac{1}{2} \sum_{i=1}^{n} \left\{ \frac{\Psi(\eta, \gamma; Y_i \mid X_i) - \Psi_0(Y_i \mid X_i)}{\Psi_0(Y_i \mid X_i)} \right\}^2 \\ & + \frac{1}{3} \sum_{i=1}^{n} t_i \left\{ \frac{\Psi(\eta, \gamma; Y_i \mid X_i) - \Psi_0(Y_i \mid X_i)}{\Psi_0(Y_i \mid X_i)} \right\}^3 \end{aligned}$$

for $\eta = C/\sqrt{n}$, where $|t_i| \leq 1$. By Taylor expansion again for $\Psi(\eta, \gamma; Y \mid X)$ at $\eta = 0$, we have $\Psi(\eta, \gamma; Y \mid X) = \Psi_0(Y \mid X) + \eta\Psi'(0, \gamma; Y \mid X) + \frac{\eta^2}{2}\Psi''(\tilde{\theta}, \gamma; Y \mid X)$, for a $\tilde{\theta} \leq \theta$. Then, by conditions C1-C5, we have

$$S_1 = \left[\sum_{i=1}^{n} \eta \frac{\Psi'(0, \gamma; Y_i \mid X_i)}{\Psi_0(Y_i \mid X_i)} - \frac{1}{2} \sum_{i=1}^{n} \eta^2 \left\{ \frac{\Psi'(0, \gamma; Y_i \mid X_i)}{\Psi_0(Y_i \mid X_i)} \right\}^2 \right] (1 + o_p(1)).$$

By Lemma 8.1 for the class $\mathcal{D}$, we know that $\frac{1}{\sqrt{n}} \sum_{i=1}^{n} \frac{\Psi'(0, \gamma; Y_i|X_i)}{\Psi_0(Y_i|X_i)}$ converges uniformly in distribution to a Gaussian process and $\sum_{i=1}^{n} \left\{ \frac{\Psi'(0, \gamma; Y_i|X_i)}{\Psi_0(Y_i|X_i)} \right\}^2 = O_p(n)$ by the law of large

numbers. Therefore,

$$S_1 = \frac{C}{\sqrt{n}} O_p(\sqrt{n}) - \frac{C^2}{n} O_p(n).$$

For S_2, we know that $|\pi_l - \pi_{0(l-K+K_0)}| = |\eta \rho_{l-K+K_0}| \leq \frac{C}{\sqrt{n}}, \quad l = K - K_0 + 1, \ldots, K$, by the restriction condition on ρ_k, $k = 1, \ldots, K_0$. Thus, by Taylor expansion, we have

$$|S_2| = \left| n\lambda \sum_{l=K-K_0+1}^{K} \frac{\pi_l - \pi_{0(l-K+K_0)}}{\epsilon + \pi_{0(l-K+K_0)}} \{1 + o(1)\} \right| = O(\sqrt{n}) \frac{CK_0}{\sqrt{n}} \{1 + o(1)\} = O(C),$$

if $\sqrt{n}\lambda \to a$. Therefore, when C is large enough, the second term in S_1 dominates S_2 and other terms in S_1. Consequently, we have $Q_{\mathrm{P}}(\eta, \gamma) - Q_{\mathrm{P}}(0, \gamma) < 0$ with probability tending to one. Hence, there exists a maximizer (η, γ) such that $\eta = O_p(n^{-1/2})$ with probability tending to one.

Then, we show that $\widehat{K} = K_0$ or $\widehat{\pi}_k = 0, k = 1, \ldots, K - K_0$ when the maximizer (η, γ) satisfies $\eta = O_p(n^{-1/2})$. We first show that, for any maximizer $Q_{\mathrm{P}}(\eta^*, \gamma^*)$ with $|\eta^*| \leq Cn^{-1/2}$, if there is a $k \leq K - K_0$ such that $Cn^{-1/2} \geq \pi_k^* > n^{-1/2}/\log n$, there exists another maximizer of $Q_{\mathrm{P}}(\eta, \gamma)$ in the area of $|\eta| \leq Cn^{-1/2}$. It is equivalent to show that $Q_{\mathrm{P}}(\eta^*, \gamma^*) < Q_{\mathrm{P}}(0, \gamma^*)$ holds with probability tending to one for any such kind maximizer $Q_{\mathrm{P}}(\eta^*, \gamma^*)$ with $|\eta^*| \leq Cn^{-1/2}$. For any $k < K - K_0 + 1$, we have

$$\begin{aligned} & Q_{\mathrm{P}}(\eta^*, \gamma^*) - Q_{\mathrm{P}}(0, \gamma^*) \leq \sum_{i=1}^{n} \{\log \Psi(\eta^*, \gamma^*; Y_i \mid X_i) - \log \Psi_0(Y_i \mid X_i)\} \\ & -n\lambda \sum_{l=K-K_0+1}^{K} \{\log(\epsilon + \pi_l^*) - \log(\epsilon + \pi_{0(l-K+K_0)})\} - n\lambda\{\log(\epsilon + \pi_k^*) - \log \epsilon\} \\ \equiv \quad & S_1 + S_2 + S_3. \end{aligned}$$

As shown before, we have $S_1 + S_2 = O_p(C^2)$. For S_3, because $\epsilon = o(n^{-1/2}/\log n)$ and $\pi_k < n^{-1/2}/\log n$, we have

$$|S_3| = O(n \cdot \frac{C}{\sqrt{n}}) \log \frac{\pi_k^*}{\epsilon} = O(n^{1/2}),$$

which implies that S_3 dominates S_1 and S_2, and hence $Q_{\mathrm{P}}(\eta^*, \gamma^*) < Q_{\mathrm{P}}(0, \gamma^*)$. So, in the following step, we only need to consider the maximizer $Q_{\mathrm{P}}(\widehat{\eta}, \widehat{\gamma})$ with $|\widehat{\eta}| \leq Cn^{-1/2}$ and $\widehat{\pi}_k < n^{-1/2}/\log n$ for $k \leq K - K_0$.

Let $Q^*(\theta) = Q_{\mathrm{P}}(\theta) - \xi(\sum_{k=1}^{K} \pi_k - 1)$, where ξ is a Lagrange multiplier. Then, it is sufficient to show that, for the maximizer (η, γ),

$$\frac{\partial Q^*(\theta)}{\partial \widehat{\pi}_k} < 0 \quad \text{for} \quad \widehat{\pi}_k < \frac{1}{\sqrt{n}\log n}, k \leq K - K_0 \tag{8.2}$$

with probability tending to one. For $k = 1, \ldots, K$, note that $\widehat{\pi}_k$ satisfies

$$\frac{\partial Q^*(\theta)}{\partial \widehat{\pi}_k} = \sum_{i=1}^{n} \frac{f_k(\beta_k; Y_i \mid X_i)}{\sum_{l=1}^{K} \widehat{\pi}_l f_l(\beta_l; Y_i \mid X_i)} - n\lambda \frac{1}{\epsilon + \widehat{\pi}_k} - \xi = 0, \tag{8.3}$$

where $f_l(\beta_l; Y_i \mid X_i) = \exp\left\{\sum_{j=1}^{m_i} q(g(X_{ij}^T \beta_l); Y_{ij})\right\}$. By the law of large numbers, the first term of (8.3) is of order $O_p(n)$. If $k > K - K_0$ and $\eta = O_p(n^{-1/2})$, we have that

$$\widehat{\pi}_k = \pi_{0(k-K+K_0)} + O_p(n^{-1/2}) > \frac{1}{2}\min\{\pi_{01}, \ldots, \pi_{0K_0}\}.$$

Hence, the second term of (8.3) is of order $O_p(n\lambda) = o_p(n)$. Thus, $\xi = O_p(n)$. If $k \leq K - K_0$, since $\widehat{\pi}_k = O_p(n^{-1/2}/\log n)$, $\lambda = a/\sqrt{n}$ and $\epsilon = o(n^{-1/2}/\log n)$, we have

$$\left\{ n\lambda \frac{1}{\epsilon + \widehat{\pi}_k} \right\} / n = \lambda \frac{1}{\epsilon + \widehat{\pi}_k} = O_p(\lambda \cdot n^{1/2} \log n) \to \infty$$

with probability tending to one. Hence, the second term in (8.3) dominates the first and third terms when $k \leq K - K_0$ and $\widehat{\pi}_k < n^{-1/2}/\log n$, which implies that (8.2) holds or equivalently, $\widehat{\pi}_k = 0$, $k = 1, \ldots, K - K_0$ with probability tending to one. This completes the proof of Theorem 2.

8.2 B. Tables

See Tables 4 and 5.

Table 4: The median and the mean squared error (MSE) on parameter estimation based on 200 replicates in Example 2. The values of MSE are times 100. For the proposed PQL and PQL2 methods, the results above are summarized based on the models with correctly specified K_0 in 200 replications.

	$\rho = 0.4$						$\rho = 0.8$					
	PQL		PQL2		QIFC		PQL		PQL2		QIFC	
	median	MSE	median	MSE	median	MSE	median	MSE	median	MSE	median	MSE
β_{10}	1.993	0.584	1.994	0.526	1.989	0.555	1.995	0.728	1.993	0.488	1.990	0.539
β_{11}	1.007	0.799	1.000	0.635	0.997	0.662	0.996	0.903	0.997	0.218	1.000	0.269
β_{12}	-1.009	1.175	-0.996	0.887	-0.996	0.913	-1.012	1.389	-1.004	0.326	-0.998	0.423
β_{13}	1.505	0.932	1.505	0.703	1.504	0.785	1.507	0.931	1.502	0.248	1.502	0.306
β_{14}	1.000	0.484	1.001	0.366	0.997	0.404	0.997	0.389	1.000	0.111	1.002	0.127
β_{20}	-4.000	0.292	-4.000	0.273	-4.003	0.284	-3.993	0.560	-3.993	0.321	-3.998	0.344
β_{21}	2.011	0.494	2.003	0.380	2.004	0.417	1.990	0.636	1.990	0.182	1.995	0.177
β_{22}	0.992	0.773	0.998	0.662	0.998	0.697	1.005	0.697	1.001	0.177	0.997	0.224
β_{23}	-2.003	0.532	-2.003	0.468	-2.001	0.473	-1.988	0.557	-1.993	0.159	-1.991	0.179
β_{24}	0.004	0.269	0.006	0.222	0.004	0.234	0.005	0.271	0.002	0.074	-0.002	0.083
β_{30}	-1.999	0.069	-2.000	0.060	-2.000	0.064	-1.994	0.112	-1.994	0.074	-1.994	0.075
β_{31}	-2.000	0.142	-1.998	0.103	-1.997	0.106	-2.000	0.129	-2.000	0.032	-2.001	0.042
β_{32}	0.995	0.198	0.996	0.137	1.002	0.150	0.993	0.201	1.000	0.053	1.001	0.070
β_{33}	0.002	0.161	0.003	0.116	0.003	0.125	-0.003	0.127	-0.002	0.033	-0.001	0.043
β_{34}	1.000	0.060	0.999	0.046	0.996	0.052	0.998	0.076	0.998	0.019	0.998	0.022
σ_1^2	0.493	0.164	0.494	0.164	0.495	0.164	0.487	0.348	0.489	0.348	0.489	0.343
σ_2^2	0.298	0.057	0.297	0.057	0.297	0.055	0.294	0.135	0.294	0.132	0.295	0.130
σ_3^2	0.099	0.005	0.099	0.005	0.100	0.005	0.097	0.010	0.097	0.010	0.097	0.010
π_1	0.299	0.002	0.299	0.002	–	–	0.291	0.008	0.291	0.008	–	–
π_2	0.301	0.002	0.301	0.002	–	–	0.296	0.004	0.296	0.004	–	–
π_3	0.400	0.007	0.400	0.007	–	–	0.413	0.019	0.413	0.019	–	–

Table 5: The median and the mean squared error (MSE) on parameter estimation based on 200 replicates in Example 3. The values of MSE are times 100. For the proposed PQL and PQL2 methods, the results above are summarized based on the models with correctly specified K_0 in 200 replications.

	PQL		PQL2		QIFC	
	median	MSE	median	MSE	median	MSE
β_{10}	1.999	0.402	1.999	0.364	1.999	0.358
β_{11}	1.002	0.596	0.994	0.305	0.998	0.341
β_{12}	-0.994	0.919	-0.996	0.494	-0.998	0.460
β_{13}	1.494	0.697	1.499	0.358	1.495	0.385
β_{14}	1.007	0.319	1.003	0.169	1.003	0.193
β_{20}	-3.988	0.280	-3.990	0.177	-3.999	0.172
β_{21}	1.993	0.419	1.993	0.223	1.999	0.218
β_{22}	1.008	0.651	1.004	0.287	1.008	0.300
β_{23}	-2.005	0.420	-2.004	0.213	-2.005	0.234
β_{24}	-0.003	0.183	0.002	0.084	0.001	0.091
β_{30}	-1.994	0.136	-1.996	0.119	-1.996	0.129
β_{31}	-1.995	0.225	-1.994	0.188	-1.993	0.205
β_{32}	0.990	0.368	0.994	0.284	0.994	0.304
β_{33}	0.001	0.210	0.003	0.156	0.005	0.195
β_{34}	0.998	0.119	0.997	0.107	0.996	0.111
β_{40}	-0.011	0.206	-0.012	0.200	-0.008	0.193
β_{41}	0.996	0.286	0.996	0.267	0.994	0.273
β_{42}	-0.010	0.438	-0.009	0.337	-0.001	0.373
β_{43}	1.013	0.369	1.008	0.307	1.008	0.301
β_{44}	1.005	0.161	1.005	0.135	1.008	0.139
β_{50}	-3.992	0.366	-3.995	0.372	-3.994	0.368
β_{51}	0.022	1.073	0.017	1.052	0.017	1.051
β_{52}	-1.010	1.266	-1.018	1.272	-1.012	1.285
β_{53}	-1.000	0.832	-1.002	0.827	-1.004	0.845
β_{54}	-1.508	0.435	-1.506	0.435	-1.509	0.435
σ_1^2	0.495	0.166	0.492	0.164	0.497	0.163
σ_2^2	0.299	0.046	0.299	0.047	0.302	0.044
σ_3^2	0.097	0.008	0.098	0.008	0.097	0.008
σ_4^2	0.142	0.021	0.143	0.019	0.146	0.017
σ_5^2	0.603	0.160	0.601	0.160	0.598	0.158
π_1	0.273	0.053	0.273	0.053	–	–
π_2	0.271	0.042	0.271	0.042	–	–
π_3	0.123	0.054	0.123	0.054	–	–
π_4	0.127	0.053	0.127	0.053	–	–
π_5	0.203	0.002	0.203	0.002	–	–